



CSC –IT Center for Science and Open OnDemand

Sebastian von Alfthan - SC23



CSC – IT Center for Science

- CSC - Non-profit company owned by the state of Finland and Finnish higher education institutions
- National supercomputers: Puhti and Mahti
 - Puhti - Supercomputer with Intel CPUs & V100 GPUs
 - Mahti - Supercomputer with AMD CPUs and A100 GPUs
- Object storage service: Allas
- LUMI
 - #5 in Top500, 379 Pflops
 - Funded by EuroHPC JU and Lumi consortium



**Mahti
Mahti-AI**



Fast parallel storage

**Puhti
Puhti-AI**




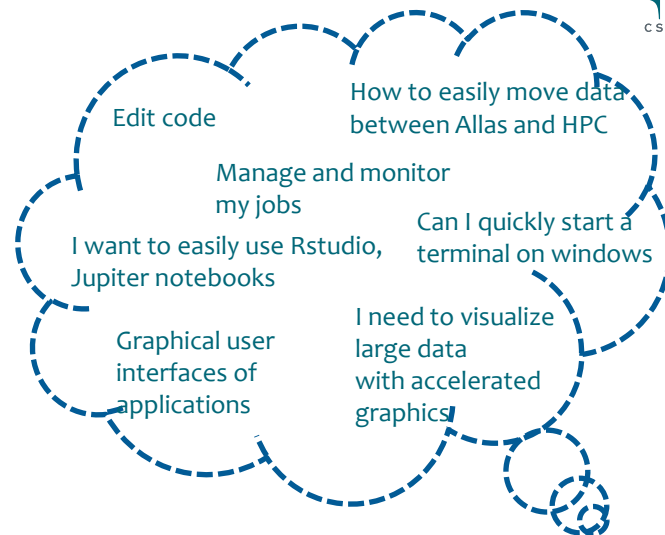
Fast parallel storage



Allas

Background

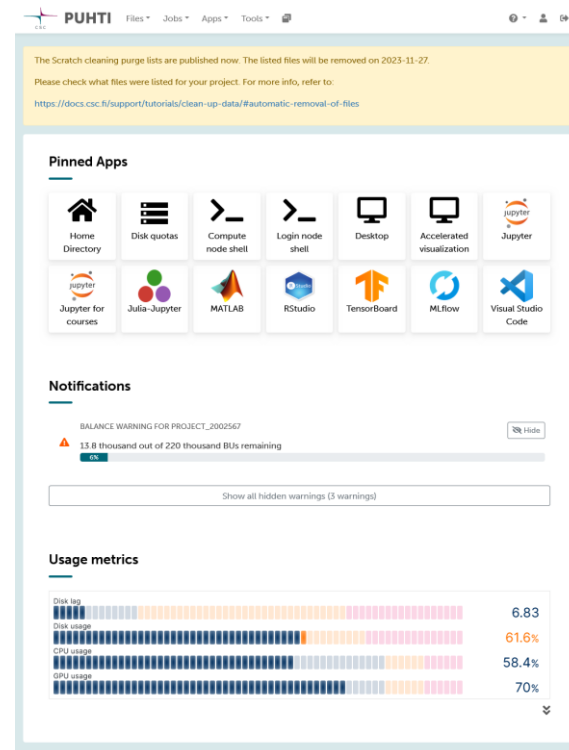
- Puhti taken into use 2019
 - Large user community – many with limited HPC background
 - Used extensively for trainings arranged by CSC and Universities
- Main driver for adding www interfaces
 1. Ease of use for non-expert users
 2. Platform for training – fast setup and customized environments
 3. New capabilities and productivity for power users



Web interfaces

- CSC has introduced web interfaces to its supercomputers
 - Puhti in October 2021
 - Mahti in June 2023
 - Lumi in November 2023
- Deployments are done based on OSC's Open OnDemand
- Principles
 - **Usability**– only offer well functioning features and streamlined UX, help users to use resources efficiently
 - **Robust operations** - CI/CD and deployment to manage, security focus
 - **Develop** and contribute upstream key features

- In production since 2021-09
 - Updated in 17 releases
- Applications
 - Jupyter notebooks – also julia and custom environments
 - Desktop environments with gpu support
 - Matlab, Rstudio, Vscod, persistent compute node ssh terminal
 - AI Apps: MLFlow, Tensorboards, pytorch notebook
- Custom
 - Passenger app for project & resource views
 - Quota notification widgets




The screenshot shows the PUHTI web interface. At the top, there is a navigation bar with 'PUHTI' and menu items like 'Files', 'Jobs', 'Apps', and 'Tools'. Below the navigation bar, there is a yellow notification banner about scratch cleaning purge lists. The main content area is divided into three sections: 'Pinned Apps', 'Notifications', and 'Usage metrics'. The 'Pinned Apps' section displays a grid of application icons including Home Directory, Disk quotas, Compute node shell, Login node shell, Desktop, Accelerated visualization, Jupyter, Jupyter for courses, Julia-Jupyter, MATLAB, RStudio, TensorBoard, MLflow, and Visual Studio Code. The 'Notifications' section shows a warning for 'PROJECT_2002567' with 13.8 thousand out of 220 thousand BUs remaining. The 'Usage metrics' section displays four horizontal bar charts for Disk I/O, Disk usage, CPU usage, and GPU usage, with values 6.83, 61.6%, 58.4%, and 70% respectively.


www.puhti.csc.fi statistics

- Steady growth in user numbers, approaching 50% of all users

Unique users of www.puhti.csc.fi



Fraction of Puhti users using www.Puhti



The Queen of North

LUMI

European flagship
supercomputer



www.lumi-supercomputer.eu

#lumisupercomputer #lumieurohpc

LUMI is an HPE Cray EX Supercomputer

LUMI




Hewlett Packard
Enterprise


LUMI is one of the fastest supercomputers in the world

LUMI


SUSTAINED PERFORMANCE

375 PETAFL0P/S

= performs 375×10^{15} calculations per second



2 x



High-performance computing

AI

Data analytics

Modern architecture

LUMI-C:
x86 Partition
Supplementary CPU partition:
200,000
AMD EPYC CPU cores.

LUMI-K:
Container Cloud Service

LUMI-O:
Object Storage Service
30 PB
encrypted object storage
(Ceph) for storing, sharing
and staging data.

LUMI-Q:
Quantum Computing

High-speed interconnect

Possibility for combining
different resources within
a single run. HPE
Slingshot technology.

LUMI-G:
GPU Partition
Sustained performance
375
Pflop/s powered by AMD
Radeon Instinct™ MI250X GPUs.

LUMI-D:
Data Analytics Partition
Interactive partition with
32 TB
of memory and graphics GPUs for
data analytics and visualization.

LUMI-F:
Accelerated Storage
10 PB
Flash-based storage layer with
extreme I/O bandwidth of
2 TB/s and IOPS capability.

LUMI-P:
Lustre Storage
80 PB
parallel file system.

www.LUMI.csc.fi deployment

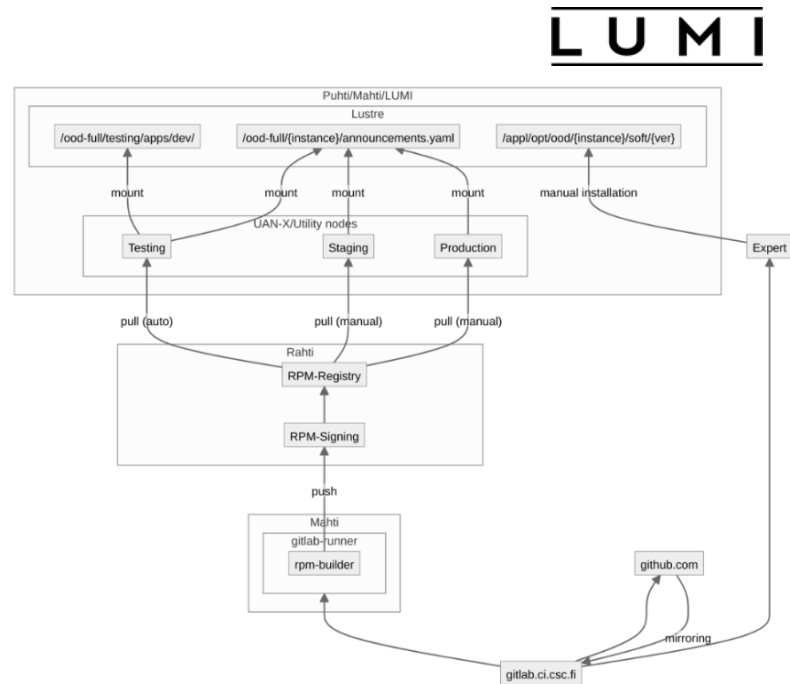
- **Lumi: November 8 2023**
- Based on Open onDemand 3.0
- Second generation of CSC OoD deployments
 - Based on Mahti deployment introduced in summer 2023
- Container + RPM based deployment on user access nodes
- Live usage statistics - Matomo

The screenshot displays the LUMI web interface. At the top, there is a navigation bar with the LUMI logo and menu items: Files, Jobs, Apps, Tools, and a search icon. On the right side of the navigation bar, there are system status icons (Wi-Fi, battery, signal) and user profile icons. The main content area is divided into three sections:

- Pinned Apps:** A grid of ten application tiles. The first row includes Home Directory, Compute node shell, Login node shell, Desktop, and Active Jobs. The second row includes Jupyter, Jupyter for courses, Julia-Jupyter, TensorBoard, and Visual Studio Code.
- Notifications:** A section titled "Notifications" with the message "You have no notifications."
- Usage metrics:** Two gauge charts. The first chart, titled "CPUs allocated", shows a value of 194445. The second chart, titled "GPUs allocated", shows a value of 18659.

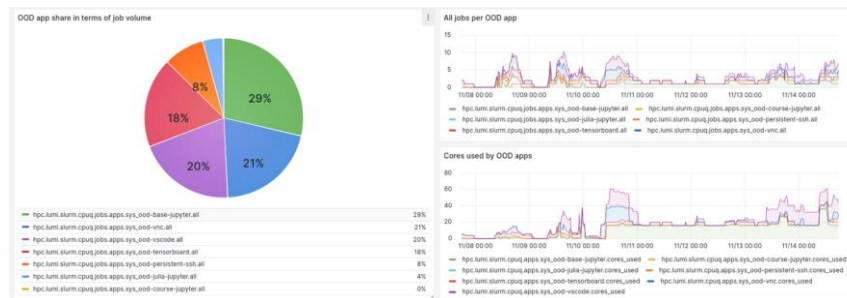
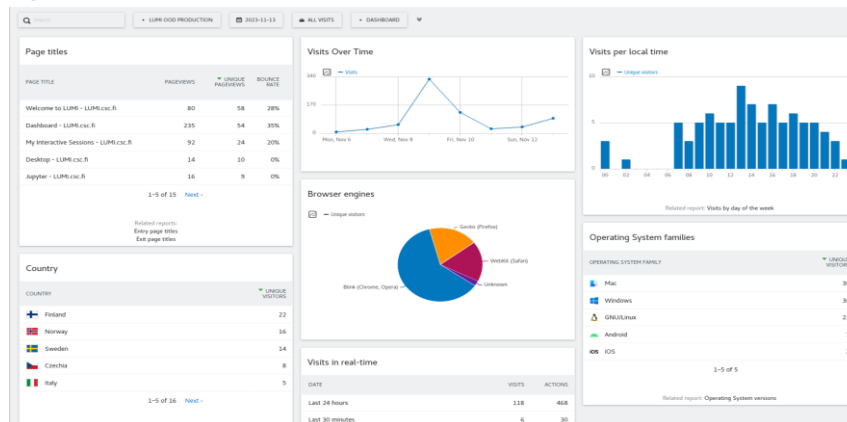
Deployment

- Production, staging and dev instances
 - RPM repo in Rahti (CSC K8s)
 - Stores app RPMs for Puhti, Mahti and LUMI OOD
 - Easy to specify apps and versions in each OOD release
 - RPMs built by GitLab CI Runner on supercomputer
- Admins build OOD container with upstream RPMs and CSC app RPMs
 - Testing instance updated every 5 minutes with latest app RPMs
 - Easy and reproducible app deployment
 - Also capability to build custom OOD RPMs



Live statistics for usage monitoring

- Website analytics based on matomo
- Selfhosted → no data leaves the HPC site
- Heavily anonymized, so some loss accuracy, but maintains user privacy and does not require any cookies.
- Grafana graphs for application distribution and currently used resources and job volumes.



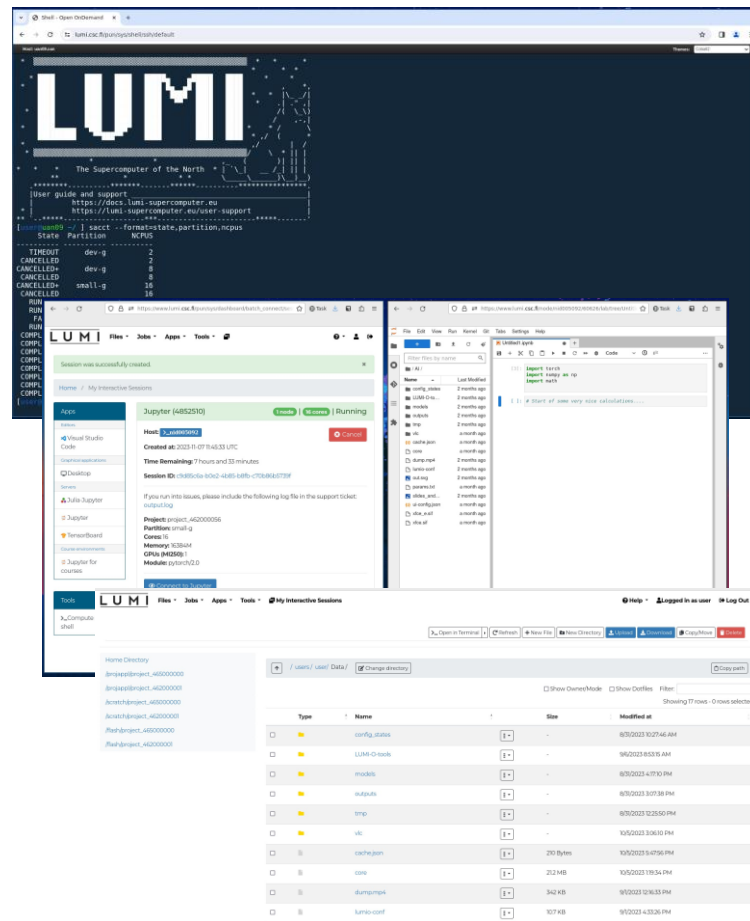
www.LUMI.csc.fi user environment

• Applications

- Jupyter notebooks – also julia and custom environments
- Desktop environments with gpu support
- Vscode, persistent compute node ssh terminal
- AI Apps: Tensorboard, pytorch notebook

• Highlights

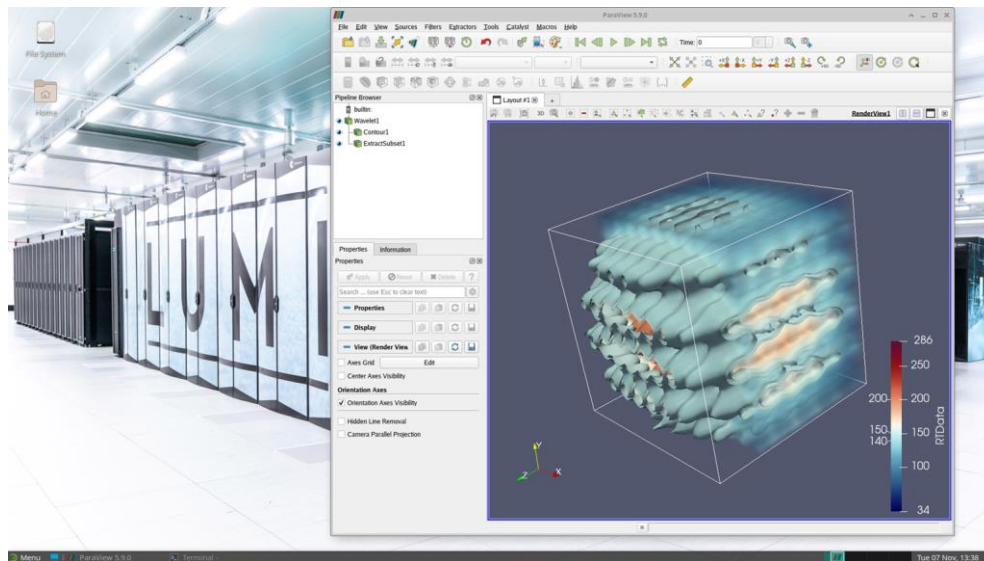
- Full production use of Lumi-D partition – interactive analytics in large memory nodes and accelerated viz
- Access to quantum computers for quantum research projects
- Lumi-O access in File manager



Accelerated graphics

LUMI

- GPU accelerated applications for visualization.
 - E.g Paraview, Blender
- Desktop and applications fully containerized.
 - Base VGL EGL container publicly available → users can use this to easily build their software offsite and the run it through the web interface with acceleration
 - ghcr.io/lumi-supercomputer/vgl:1.0



Specialized environments

- Helmi 5-qubit quantum computer co-developed by VTT and IQM accessible from LUMI
- Now enabled notebooks to run code on quantum computer
- Looking at creating more domain specific profiles for AI, Quantum, ... to cater to specific needs

```
[1]: import os
from qiskit import QuantumCircuit, execute
from qiskit_iqm import IQMProvider

HELMI_CORTEX_URL = os.getenv('HELMI_CORTEX_URL') # This is set when loading the module

provider = IQMProvider(HELMI_CORTEX_URL)
backend = provider.get_backend()

shots = 1000 # Set the number of shots you wish to run with

# Create your quantum circuit.
# Here is an example
circuit = QuantumCircuit(2, 2)
circuit.h(0)
circuit.cx(0, 1)
circuit.measure_all()

print(circuit.draw(output='text'))

job = execute(circuit, backend, shots=shots) # execute your quantum circuit
counts = job.result().get_counts()
print(counts)
```

q_0: H ———— M

q_1: ———— X ———— M

c: 2/

meas: 2/

0 1

{'10': 100, '00': 451, '11': 388, '01': 61}

[]:

Cloud storage integration



- We wanted LUMI-O S3 object storage to be usable through OOD
- Existing solutions were not user-friendly or possible to use
 - Mounts not possible
 - File manager as interactive app works, but inconvenient
- Decided to develop cloud storage support for the OOD files app
- Core implementation submitted in PRs during 2022
 - Many thanks to OSC for PR reviews
- Cloud storage support released in OOD in 3.0

Cloud storage integration: Allas and Lumi-O



- Uses Rclone to support >40 different storage systems
 - S3, Swift, Microsoft OneDrive, Google Drive, etc.
- Rclone configuration read from user's home directory
 - Remotes are validated and added as shortcuts (favorite paths)
 - Configuration through SSH or web UIs
- Supports almost all actions supported by OOD files app
 - Directory download as ZIP not supported
 - Uploading large files is not possible

The screenshot shows the MAHTI web interface for cloud storage configuration. At the top, there is a navigation bar with 'MAHTI' and several menu items: 'Files', 'Jobs', 'Apps', 'Tools', and a user profile icon. Below the navigation bar, a green banner states: 'The web interface has been updated to release 2. Visit <https://docs.csc.fi/support/wn/comp-new/> to see what's new.'

Cloud storage configuration

Currently configured remotes:

Name	Type	Expiry	Action
s3allas-project_2001316	s3		Delete Revoke
s3allas-project_2001659	s3		Delete Revoke

The table above lists the currently configured remotes in the Rclone config.

Deleting a remote only removes the configured Rclone remote, while keeping the remote's access token valid as it may be used outside of the web interface. To revoke the access token too, you can use the revoke button for the remote, which will revoke the access token and delete the remote.

New Rclone remotes for accessing Allas can be created below. After configuring a remote, you will be able to access it in the Files dropdown in the navigation menu. When transferring files in the file browser, keep the browser tab with the transfer open to ensure that it completes successfully.

Configure new remotes:

Allas S3 Allas Swift

If an S3 access token already exists for the project, it will be reused for the remote. Otherwise, a new S3 access token will be generated. The S3 access tokens are valid forever.

Project

project_2001990

Submit

Future outlook

- Now focused on interactive use – looking at improved job submission capabilities and workflow tools with web UI
- Developing improved backend for supporting large data transfers
- Leveraging domain specific profiles and new features in 3.1
- Releasing more recipes, apps
 - <https://github.com/search?q=org%3ACSCfi%20ood&type=repositories>